Whitepaper

AI Powered Pricing Strategies –

Predictive Models and Elasticity Analysis

Table of **Contents**

Executive Summary

This whitepaper discusses the importance of product pricing in business, highlighting value-based pricing. It emphasizes that the pricing process, which is crucial for revenue, profits, reputation, and an organization's future, is often manual, time-consuming, and requires technical expertise. The paper introduces an AI-based solution for automatic price prediction and price elasticity analysis to optimize the pricing process. This advanced AI-based pricing analytics solution enables businesses to determine the optimal price point that balances demand and value, thereby automating the pricing mechanism. The paper explores the benefits and drawbacks of this AI-based solution, providing valuable insights for businesses looking to enhance their pricing strategies.

Whether you're a business owner, demand planner, or operations manager looking to automate pricing through AI, or if you want to understand which product features impact pricing, this paper offers valuable insights and solutions tailored to your requirements.

Introduction

Pricing products appropriately is critical for success. Set prices too low, and potential earnings are forfeited; set them too high, and customers may turn to competitors. Striking a balance between profitability and customer satisfaction is therefore paramount.

This whitepaper addresses the challenge of determining the optimal price point, which involves predicting product prices based on specifications, quantity, and temporal features, and identifying the price elasticity — the point at which pricing becomes most effective.

We introduce an AI-based solution that automates price prediction, extracts key performance indicators (price drivers), enables real-time price predictions, and justifies the predicted price using explainable AI. Furthermore, it determines the optimal price point through price elasticity analysis.

Businesses can automate and optimize their pricing strategies by leveraging advanced AI techniques, ensuring they remain competitive while maximizing profitability. This whitepaper provides a comprehensive overview of this AI-based solution, its benefits, and its potential applications in various business contexts.

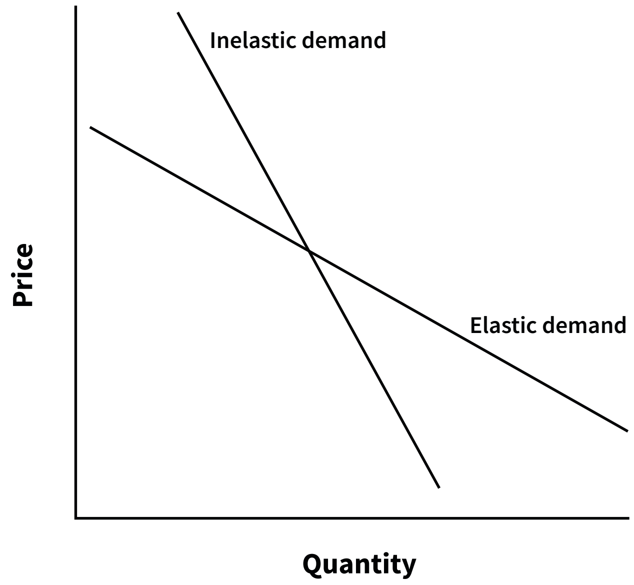

Solution

Our solution employs a streamlined process, as depicted in the provided flowchart. Initially, the user inputs product specifications into the system. The solution then enhances this data through feature engineering to generate additional attributes. Subsequently, two predictive models are utilized: a Price Prediction model and a Price Elasticity model. The Price Prediction model forecasts the product's base price, while the Price Elasticity model determines the most profitable price point. This dual-model approach ensures an optimal balance between accurate price prediction and profit maximization. The final output is the optimal price, which results from the combined predictions of both models. This comprehensive process allows us to leverage user inputs and advanced modeling techniques to deliver a robust pricing solution.

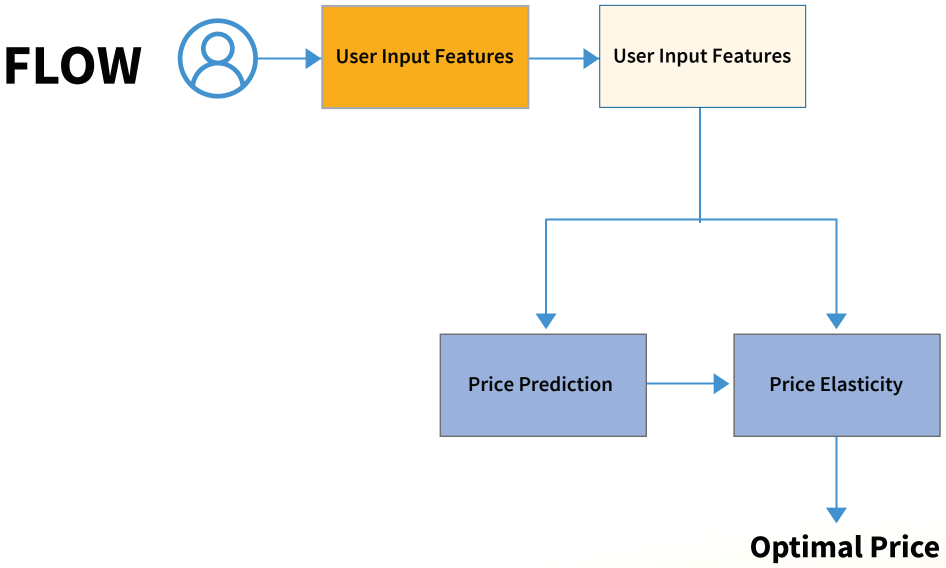

Price Prediction:

This document presents a comprehensive guide to developing a machine-learning model for product price prediction. The model predicts product's initial price in real-time based on historical data, and a limited set of price-driving features.

The data required for this model includes product attributes (length, width, color, weight, shape, etc.), temporal features (year, month, date, etc.), hierarchical features (country, state, city, customer, product type, etc.), order quantity, features engineered from the data, and external factors (market trends, economic indicators, competitor pricing).

The process of price prediction presents several challenges. These include feature complexity, where not all features directly impact the price, and some features only contribute when considered with others. Data quality and anomalies can also pose challenges, with records containing unusual values due to manual errors or data entry mistakes. Outliers can adversely affect model accuracy, and missing data needs to be addressed using strategies like imputation or removal. Another challenge is handling large volumes of data efficiently during preprocessing, modeling, and evaluation.

Several data transformations are necessary to address these challenges. These include data cleaning and preprocessing, outlier detection and handling, handling categorical variables, feature engineering, and feature scaling.

Data cleaning and preprocessing involve removing special characters and extra spaces, converting text to lowercase, dropping columns with null values exceeding a certain threshold, filtering out records with unusual or unrealistic prices, eliminating features with zero variance, and addressing multicollinearity.

Outlier detection and handling involve techniques like the Isolation Forest algorithm to identify and remove outliers. Handling categorical variables involves dealing with rare or infrequent categorical values by grouping or encoding them differently. Feature engineering involves binning numeric features into buckets when the exact value is not critical during model inference. Feature scaling involves standardizing numeric features using techniques like Standard Scaler.

These steps are essential for preparing data before building any machine learning model, regardless of the specific domain or application.

Price Elasticity:

Our pricing strategy employs two models. The first predicts base prices based on historical data, while the second is necessary for for-profit maximization, classification of the price elasticity to pinpoint the optimal price point.

Price Elasticity of Demand (PED) gauges the sensitivity of demand to price changes. Defined as the ratio of the percentage change in demand to the percentage change in price (PED = %Δdemand / %Δ price), it helps understand market dynamics.

A PED value ≥1 indicates elastic demand, where consumers are highly responsive to price changes, leading to a significant decrease in demand for minor price hikes. Conversely, a PED value <1 signifies inelastic demand, where consumers are less sensitive to price changes, resulting in a smaller decrease in demand for price increases.

We conduct this analysis at the SKU level, considering the hierarchy: Country > Customer > Product (SKU). The data includes country, customer, product, quantity, and price. For instance, to predict elasticity at the country-customer-product level, we need data in this format:

Table				
Country	Customer	Product	Quantity	Price
China	John Doe LLC	LABEL	25000	$34

With the above data, it will be difficult to get the models which can accurately find the optimal price point, we need few extra features which indicate the demand patterns, customer relationships and geographic limitations etc.

Demand classification:

Any demand is classified into four Types: "Smooth," "Intermittent," "Lumpy," and "Erratic." These are Based on two statistical metrics: "Average Demand Interval" and "Square of Coefficient of Variation." The variability in demand timing can be analyzed using the average demand interval, and the variability in demand quantity can be analyzed using the square of the coefficient of variation.

The Average Demand Interval (ADI). It measures the demand regularity in time by computing the average interval between two demands.

- ADI = Total number of periods / Number of demand Buckets

  The square of the Coefficient of Variation (CV^2). It measures the variation in quantities.

- Coefficient of Variation = Standard deviation of population / average value of population

- CV2 = Coefficient of Variation * Coefficient of Variation

The demand classification is extracted for all the unique SKUs based on ADI and CV2.

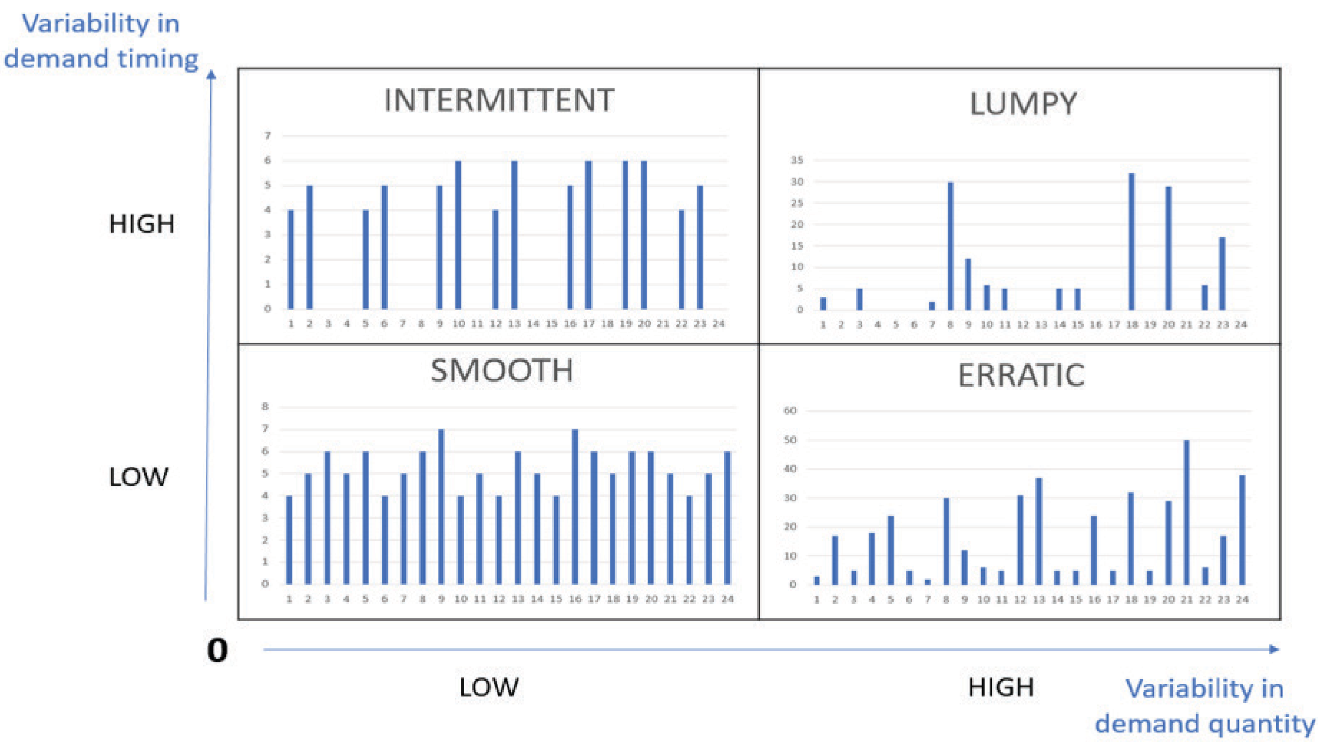

Pareto analysis:

As per Pareto analysis, Pareto analysis is premised on the idea that 80% of revenue can be achieved by doing 20% of the SKUs or, conversely, 80% of problems can be traced to 20% of the causes. Pareto analysis is a powerful quality and decision-making tool. In the most general sense, it is a technique for getting the necessary facts needed for setting priorities.

These customers' characteristics are generally different from others, these customers are giving orders throughout the year in big numbers and helping the business to sustain. The classification is divided into three categories A, B and C.

- **Category A:** These high-value items contribute significantly to the inventory value. Although they may represent a smaller proportion of the total item count, they substantially impact the business's financial performance.
- **Category B:** These items have moderate value and significance. They occupy an intermediate position between category A and C items regarding their contribution to the inventory value and item count.
- **Category C:** These items have relatively low value and contribute to a smaller portion of the overall inventory value. However, they often constitute most of the item count.

The following are the threshold values used for classification.

Class	Number of Items	Revenue
A	5%	40%
B	15%	40%
C	80%	20%

The Pareto analysis can be done at multiple levels like country, product or customer. Once the new features are ready, we can split the data into train and test and train the machine learning models.

The process of modelling involves several key steps:

The methodologies employed in both Price Prediction and Price Elasticity modelling share a common foundation, albeit with a key distinction. Price Prediction utilizes a regression model, aiming to predict the numerical value of a product's price. Conversely, Price Elasticity employs a classification model, to determine the category of elasticity, either Elastic or Inelastic. This fundamental difference like the models – numeric prediction versus categorical classification – differentiates these two approaches, despite their similar process structure.

- **Algorithm selection:** Regression models are trained using algorithms such as Random Forest, XGBoost, and other ensemble methods. The choice of algorithm depends on its suitability for the problem and dataset.

- **Hyperparameter Tuning:** Hyperparameters, which dictate the behaviour of the machine learning algorithm, are automatically tuned using the OPTUNA framework. This framework leverages Bayesian Optimization to explore different hyperparameter combinations and find the optimal configuration.

- **Error Computation and Model Refinement:** The model's performance is evaluated using an error metric after each trial. Based on this error, OPTUNA adjusts the hyperparameters and retrains the model in an iterative process until the final trial.

- **Workflow Management with MLFlow:** The entire machine learning workflow is managed using tools like MLFlow, which helps manage experiments, track metrics, and store models. It also facilitates the comparison of different model versions and collaboration among team members.

- **Champion Model and Model Registry:** The best-performing model, known as the "champion" model, is identified through hyperparameter tuning. This model is registered in the Model Registry for easy tracking and deployment. In serverless model inferencing, the champion model is used in production environments.

Model interpretations:

Now the models are trained, deployed and ready for real time inferencing but how do the customer trust these models. How do the customer believe that the model predictions are making sense or not.

We can achieve this using model interpretation, data scientists must ensure that model predictions are visible, and stakeholders accept the model output, even in cases where the model performs well. With explainable AI, we can decipher intricate Blackbox machine learning algorithms.

There are two ways of interpreting these models:

- **Local interpretations**
  - Interpreting every predicted value using algorithms like LIME or SHAP.
  - SHAP is a framework that uses Shapley values, a game theory method, to explain model outputs. It can be applied to any model but is more efficient with specific types, like tree ensembles. SHAP values, due to their additive nature, can be used for both local and global explanations and serve as a basis for advanced ML analysis including model monitoring, fairness, and cohort analysis.
  - Local Interpretable Model-Agnostic Explanations (LIME) is a technique that creates a transparent model around the decision space of any complex model's prediction. It focuses on modeling the local neighborhood of a prediction, allowing even simple linear models to approximate complex model behavior. LIME generates synthetic data by perturbing individual data points, which is then evaluated by the complex model and used to train the transparent model. This method allows for interpretations similar to linear models and can be applied to nearly any model, though explanations can sometimes be unstable and dependent on the perturbation process.
- **Global Interpretations**
  - Partial dependency plots

  \- Partial dependence plots are graphical representations that illustrate the correlation between the outcome and a subset of target features, typically one or two, while considering the influence of all other features. This method, based on perturbations, is relatively efficient in terms of interpretability. However, it operates under the assumption of feature independence, and can potentially provide misleading interpretations when this condition is not satisfied, such as in cases where the model includes numerous high-order interactions.

- Morris Sensitivity analysis

- The Morris method, also referred to as a one-step-at-a-time (OAT) global sensitivity analysis, is a technique where only one input variable is modified per iteration. This method is known for its efficiency (requiring fewer model executions) compared to other sensitivity analysis algorithms. However, it has the limitation of being unable distinguish between non-linearities and interactions. It is typically employed for preliminary screening to identify which inputs warrant further examination. The implementation of this method is facilitated using SALib.

- Feature Importance plots

- Extracts the key price drivers.

Business Benefits

Below are a few of the benefits of our proposed approach:

Revenue	Efforts	Time
Now, the client can sell products based on the real value of the products compared to cost-plus-based pricing.	The end user does not need to enter hundreds of variable values into the system to get the base price, as the ML model needs only a few variables to predict the price.	The customer-facing team can instantly quote the price to their customers, compared to a two-week-long tedious process.

Conclusion

This AI-based pricing solution automatically predicts the base price of the products and suggests optimal prices, providing interpretations of complex machine learning models using explainable AI.

AI-based pricing solutions offer numerous advantages including real-time price optimization, personalized pricing strategies, profit maximization, improved decision support, efficiency, scalability, market responsiveness, and enhanced revenue and market share. These benefits make AI an invaluable asset for businesses. However, of implementing such solutions can be complex and there may be an overreliance on algorithms.

Reference:

1. Hyper Parameter Tuning using OPTUNA
2. Demand pattern cgassification
3. SHAP modeg interpretation
4. Partiag dependency Pgots

About the Author

Naveen Kumar Gandla

Specialist - Data Sciences

Naveen is a committed data scientist with a fervor for developing cutting-edge AI solutions. His proficiency in gathering requirements, designing solutions, and implementing them has been honed over seven years of industry experience. He has worked across a various of sectors, including retail, banking, FMCG, CPG, and manufacturing. His dedication and expertise make him an asset data science.

About LTIMindtree

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700 clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by 81,000+ talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit **https://www.ltimindtree.com/.**